

Adaptive Controller Effects on Pilot Behavior

Anna Trujillo and Irene Gregory, Ph.D.

Research Directorate
NASA Langley Research Center
Hampton, USA

[anna.c.trujillo;irene.m.gregory]@nasa.gov

Lucas Hempley

Northrop Grumman Technical Services
Hampton, USA

Abstract—Adaptive control provides robustness and resilience for highly uncertain, and potentially unpredictable, flight dynamics characteristic. Some of the recent flight experiences of pilot-in-the-loop with an adaptive controller have exhibited unpredicted interactions. In retrospect, this is not surprising once it is realized that there are now two adaptive controllers interacting, the software adaptive control system and the pilot. An experiment was conducted to categorize these interactions on the pilot with an adaptive controller during control surface failures. One of the objectives of this experiment was to determine how the adaptation time of the controller affects pilots. The pitch and roll errors, and stick input increased for increasing adaptation time and during the segment when the adaptive controller was adapting. Not surprisingly, altitude, cross track and angle deviations, and vertical velocity also increase during the failure and then slowly return to pre-failure levels. Subjects may change their behavior even as an adaptive controller is adapting with additional stick inputs. Therefore, the adaptive controller should adapt as fast as possible to minimize flight track errors. This will minimize undesirable interactions between the pilot and the adaptive controller and maintain maneuvering precision.

Keywords—adaptive controller; adaptation time; piloting effects

I. INTRODUCTION

Adaptive control in flight applications has a long and rich history dating back to the 1950s. Adaptive control provides robustness and resilience for highly uncertain, and potentially unpredictable, flight dynamics that characterize stall conditions or are due to damage on transport as well as high-performance aircraft. In the past decade, some flight experiences of pilots with an adaptive controller have exhibited unpredicted interactions [1, 2]. In retrospect, this is not surprising once it is realized that there are now two adaptive controllers interacting – the software adaptive control system and the pilot. The pilot is controlling the attitude and rates of the vehicle (definition of a control system), and the pilot’s method of controlling will change due to dynamically changing system parameters. One hypothesized reason for the adverse interactions with an adaptive controller in the loop is the pilot not realizing how the adaptive controller is changing aircraft dynamics.

A. Previous Research

As shown in [1] and [2], the pilot-adaptive controller interaction may be unpredictable. It is unclear, though, what is causing these interactions assuming the adaptive flight controller is adapting in near real-time. In fact, others essentially found no differences in piloting capabilities and preferences

using an adaptive controller except for low dynamic pressure flight phases where pilots rated the adaptive controller as better [3]. However, past data suggested that the “gain-changing logic can be fooled by ... pilot control activity” [3].

More recent research is studying the efficacy of using adaptive controllers to alleviate inappropriate pilot-control coupling caused by rate limiting and a pure time delay [4]. The delay margin of the pilot-vehicle system is estimated in the presence of rate limiting and time delay and an \mathcal{L}_1 adaptive controller is inserted to compensate and achieve the desired response of the system [5].

With these results in mind, this research set out to determine the effect that the speed of the adapting controller would have on a pilot during a manual control task. The work would contribute toward characterizing the mechanisms by which these two adaptive controllers in the loop function and interact – one, a human and the other, a software control system.

B. Experiment Objective

An experiment was conducted to categorize the interaction on the pilot with an adaptive controller during control surface failures. The primary objective of the experiment was to determine how the adaptation time of the controller affects pilots.

II. METHOD

This experiment considered whether an adaptive controller helps pilots during control surface failures and how this controller affects pilots before, during, and after the control surface failures. The control surface failures were either a stuck surface or a slow uncommanded drift. The control surfaces failures were a combination of three control surfaces and these could be the left and right elevators symmetrically or asymmetrically, the left and right ailerons, and the upper and lower rudders. A human-in-the-loop experiment examined the effects of these control surface failures on pilot performance during a cruise flight phase while initiating a climb or descent at 500 feet per minute, or a heading change maneuver of less than 30 degrees. These maneuvers were indicated on the primary flight display (PFD) via the flight director and on the horizontal and vertical navigation map displays.

A. Simulation Environment

The physical setup of the simulator incorporated an out-the-window view in the upper center 30-inch diagonal screen and four 20-inch touchscreens below the out-the-window screen [6,

7]. The middle-left touchscreen depicted the PFD and the middle-right touchscreen depicted an engine indication display (EID). The far-right touchscreen displayed the after-run questions. Subjects flew the aircraft with a right-handed joystick. For a more detailed description of the experiment setup, see [6, 7].

B. Independent Variables

The independent variables were:

1) *Adaptation Time*: Each subject experienced four controller adaptation times: zero seconds, three seconds, seven seconds, and no adaptation (Never). These times indicated how long it took the adaptive controller to compensate for the failures and return the closed-loop aircraft dynamic response back to the desired state. Zero seconds indicated the fastest possible adaptation time, essentially the processor speed. Three seconds was chosen because with this time, the subject might notice the controller adapting. As for seven seconds, this was chosen because the subject should notice the controller adapting.

2) *Segment*: Each scenario was broken down into segments for data analyses. “Before failure” indicates the time before a failure occurred. “During failure” was while the adaptive controller was adapting. When the adaptive controller adapted immediately, there is no “during failure” segment. “After failure” was the time after the adaptive controller fully adapted to compensate for the failure. When the adaptive controller loop was open – baseline controller only – and never adapted, there was no “after failure” segment.

3) *Subjects*: The seventeen subjects were an average of 48 ± 10 years old with the youngest 29 years old and the oldest 61 years old. All of them were airline transport-rated pilots with an average of 26 ± 11 years of flight experience (minimum flight experience was 7 years and the maximum 45 years) and an average of $10,706 \pm 7164$ hours of flight experience (minimum flight hours were 2,100 and the maximum 23,400).

C. Dependent Variables

The primary dependent variables involved flight technical data. In particular in the lateral axis was: cross track error, the difference between current aircraft position and commanded position; roll; and roll error, the difference between current bank angle and commanded bank angle. In the longitudinal axis was: altitude error, the difference between current aircraft altitude and commanded altitude; pitch; vertical velocity; and pitch error, the difference between current aircraft pitch angle and commanded pitch angle. Also recorded was lateral and longitudinal stick input, and the velocity and acceleration of these stick inputs.

Two other secondary dependent variables involved subjective ratings by the participant. After each run, subjects provided a Cooper-Harper (CH) handling qualities rating (HQR) [8-10]. After certain runs, subjects also gave a NASA-Task Load Index (TLX) workload rating [6, 11].

D. Procedure

During each run, flight technical data were recorded. After each run, subjects gave a CH rating and after certain runs, a NASA-TLX workload rating. After all the data runs were

completed, subjects filled out a final questionnaire asking them about their preferences on the information in the displays and displays themselves.

III. RESULTS

A. Piloting Changes in the Longitudinal Axis

When looking at subjects’ longitudinal stick input versus time, any effects of how the adaptive controller interacts with the input are masked (Fig. 1). For example, it is unclear whether the adaptation time affected the magnitude of the longitudinal stick input during or after the failure because of the seemingly random longitudinal stick input and the amount of longitudinal stick inputs needed to maintain desired pitch command involved many fine movements every few seconds. However, when looking at the average pitch error (difference between the actual and commanded pitch) for the data runs, this increased with adaptation time ($F(3, 2913)=3.764$; $p=0.01$) and “during failure” ($F(2, 2913)=92.031$; $p \leq 0.001$) before decreasing back down to “before failure” levels after the failure (Table I). This is also reflected in the average normalized longitudinal stick input for the data runs ($F(4, 2913)=115$; $p \leq 0.001$) (Fig. 2). Hence, subjects’ longitudinal stick input and the associated pitch error increases while the failure occurs and before the adaptive controller fully adapts but they settle to pre-failure levels after the adaptive controller finishes adapting.




Fig. 1. Example Subject Normalized Longitudinal Stick Input During a Data Run

TABLE I. AVERAGE PITCH ERROR FOR THE DATA RUNS BY ADAPTATION TIME AND BY SEGMENT

Adaptation Time (seconds)	Pitch Error (degrees)	SE ^a of the Mean
0	1.24	0.018
3	1.27	0.022
7	1.30	0.022
Never	1.46	0.033
Segment		
Before Failure	1.28	0.017
During Failure	1.55	0.029
After Failure	1.12	0.013

^a. SE = Standard Error




Fig. 2. Average Normalized Longitudinal Stick Input for the Data Runs by Adaptation Time and Segment

On the other hand, the average longitudinal stick velocity ($F(4, 2913)=4.76$; $p=0.001$) (Fig. 3) and acceleration ($F(2, 2913)=74$; $p\leq 0.001$) (Table II) for the data runs slowly decrease with time. This effect may be more of a learning effect in that throughout the data run, subjects learned to better control the vehicle. Consequently, if the failure occurred later in the run, the pitch error “before failure” values may more closely mimic the pitch error “after failure” values. Two observations support this. The average longitudinal stick velocity for the data runs when the system never adapted is less than the average longitudinal stick velocity for the other adaptation times (Fig. 3). And, the longitudinal stick acceleration also decreased during the data run (Table II). In addition, while longitudinal stick velocity increased during the adaptation time but both longitudinal stick velocity and acceleration eventually settle to at or below pre-failure levels after the adaptive controller fully adapted.




Fig. 3. Average Longitudinal Stick Velocity for the Data Runs by Adaptation Time and Segment

TABLE II. AVERAGE LONGITUDINAL STICK ACCELERATION FOR THE DATA RUNS BY SEGMENT

Segment	Longitudinal Stick Acceleration	SE of the Mean
Before Failure	4.12	0.104
During Failure	3.18	0.083
After Failure	2.58	0.077

The average altitude deviation from commanded ($F(4, 293)=157$; $p\leq 0.001$) and the average vehicle’s vertical velocity ($F(4, 2913)=76$; $p\leq 0.001$) for the data runs also increased during the failure (Fig. 4); however, these values also increased after the adaptive controller fully adapted. This is true even for the 0-second adaptation time. A possible explanation for this is, given more time, the altitude deviation and vertical velocity would settle closer to “before failure” levels either due to learning effects or the fact that the longitudinal stick velocity and acceleration decrease “after failure” indicating that subjects were more closely aligned with the flight director. The second interpretation is supported by the values for these parameters during the final 20 seconds of the data run (Fig. 5).




Fig. 4. Average Altitude Deviation for the Data Runs by Adaptation Time and Segment




Fig. 5. Average Vertical Velocity for the Data Runs by Adaptation Time and Segment

B. Piloting Changes in the Lateral Axis

Similar patterns emerged for the lateral axis as there were for the longitudinal axis. As before, inspecting the lateral stick input against time does not immediately indicate any changes in subject input behavior due to the failure (Fig. 6). Once again, it is unclear whether the adaptation time affected the magnitude of the lateral stick input during or after the failure because of the seemingly random lateral stick input and the amount of lateral stick inputs needed to maintain desired roll command involved many fine movements every few seconds. However, the average roll error (the difference between commanded and actual roll) for the data runs generally increased “during failure” but decreased “after failure” ($F(4, 2913)=29$; $p<0.001$) (Fig. 7). This same pattern is reflected in the average normalized lateral stick input for the data runs ($F(4, 2913)=99$; $p<0.001$) (Fig. 8). As with longitudinal stick, subjects’ lateral stick input and the associated roll error increased while the failure occurs and before the adaptive controller fully adapted but they generally settle to pre-failure levels after the adaptive controller fully adapted. The average lateral stick input velocity ($F(4, 2913)=24$; $p<0.001$) (Fig. 9) and acceleration ($F(4, 2913)=4.756$; $p=0.001$) (Fig. 10) for the data runs also demonstrate this same pattern where the velocity and acceleration typically return to “before failure” values after the adaptive controller has fully adapted indicating nominal piloting techniques.




Fig. 6. Example Subject Normalized Lateral Stick Input




Fig. 7. Average Roll Error for the Data Runs by Adaptation Time and Segment




Fig. 8. Average Normalized Lateral Stick Input for the Data Runs by Adaptation Time and Segment




Fig. 9. Average Lateral Stick Input Velocity for the Data Runs by Adaptation Time and Segment




Fig. 10. Average Lateral Stick Input Acceleration for the Data Runs by Adaptation Time and Segment

As with altitude deviation and vertical velocity, the average cross track error ($F(4, 293)=33$; $p\leq0.001$) (Fig. 11) and bank angle error ($F(4, 2913)=51$; $p\leq0.001$) for the data runs increased during the failure (Fig. 12); however, these errors also increased after the adaptive controller fully adapted. This is true even for the 0-second adaptation time. As with the longitudinal axis, this may indicate that the vehicle dynamics felt the same to the subject but the vehicle was not exactly responding the same as it was “before failure” or potentially given more time, these errors would settle closer to “before failure” levels. The latter may be the case when looking at the values for these parameters during the final 20 seconds of the data run (Fig. 11 and Fig. 12). Hence, cross track error and bank angle error increased during the adaptation time but appear to eventually settle to pre-failure levels after the adaptive controller fully adapts.




Fig. 11. Average Cross Track Error for the Data Runs by Adaptation Time and Segment

C. Subjective Ratings

Adaptation time was significant for several of the NASA-TLX workload measures; in particular for mental demand ($F(3,764)=3.62$, $p=0.013$), temporal demand ($F(3,764)=4.55$,




Fig. 12. Average Bank Angle Error for the Data Runs by Adaptation Time and Segment

$p=0.004$), performance ($F(3,764)=4.55$, $p=0.004$), effort ($F(3,764)=5.30$, $p\leq0.01$), and overall workload ($F(3,764)=5.29$, $p\leq0.01$) (Table III and Fig. 13). As indicated in Table III, all measurements on the NASA-TLX had the same pattern with no adaptation having a higher workload than the other three adaptation times. This suggests that the adaptive controller did decrease the workload of the subject during control surface failures. The lower workload with the adaptive controller is also supported by the longitudinal and lateral stick inputs (Fig. 2 and Fig. 8, respectively) of subjects where more gross movements are present when there is no adaptive controller.

TABLE III. MEAN WORKLOAD BY ADAPTATION TIME

Workload Measure	Mean Value for Adaptation Time:			
	0 seconds	3 seconds	7 seconds	Never
Mental Demand ^{a,b}	19.85 ^c	22.32	22.38	33.10
Physical Demand	20.84	22.70	20.35	35.39
Temporal Demand ^a	20.16	23.36	22.20	34.34
Performance ^a	23.27 ^d	26.05	22.81	35.06
Effort ^a	25.94	29.65	26.39	42.90
Frustration	19.81	22.43	21.02	31.03
Overall Workload ^a	21.64	24.42	22.52	35.29

^b. * indicates statistical significance

^c. 0=low demand, effort, frustration, or overall workload and 100=high demand, effort, frustration, or overall workload

^d. 0=good performance and 100=bad performance




Fig. 13. NASA-TLX Ratings by Adaptation Time

Adaptation time was also significant for the CH rating ($X^2(3)=80.53$, $p\leq0.01$). The HQRs generally improved with the adaptive controller adaptation time (Fig. 14 and Fig. 15). As seen in the data above, having no adaptive controller increases workload (Table III and Fig. 13) and increases pitch and roll error (Table I and Fig. 7 respectively) which is reflected in the higher (*i.e.*, worse) HQRs. Conversely, the adaptive controller improved HQRs especially in the case of 0 sec adaption time, where the vast majority of ratings were between HQRs 1 and 3 (Level 1 handling qualities).

IV. CONCLUSIONS

Adaptive control provides robustness and resilience for highly uncertain, and potentially unpredictable, flight dynamics that characterize stall conditions or are due to damage on

transport as well as high-performance aircraft. But with an adaptive controller and a pilot, there are now two adaptive systems interacting. The pilot is controlling the attitude and rates of the vehicle (definition of a control system), and the pilot's method of control will change due to dynamically changing system.




Fig. 14. Count of Cooper-Harper Ratings by Adaptation Time

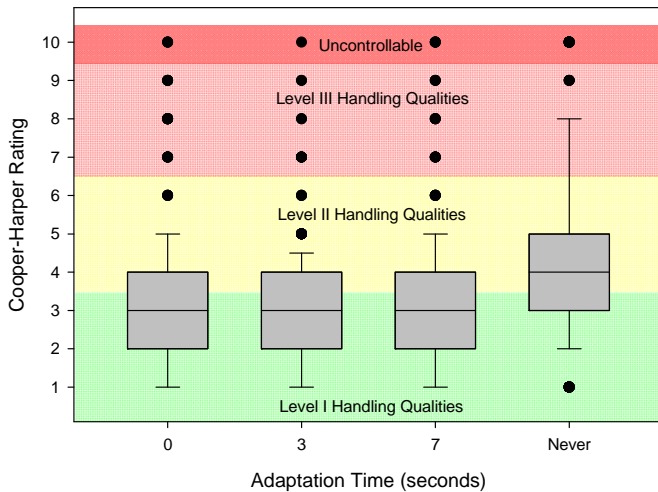


Fig. 15. Cooper-Harper Rating Box Plot by Adaptation Time

parameters. This experiment looked at how the adaptation time of the controller affects piloting inputs.

The pitch and roll errors increased for increasing adaptation time and during the segment when the adaptive controller was adapting. This is also reflected in the stick input displacement with stick input approaching pre-failure levels after the adaptive controller finishes adapting. However, stick velocity and acceleration slowly decrease with time. This may be more of a learning effect in that throughout the data run, subjects learned to better control the vehicle; however, if the failure occurred later in the run, the “before failure” values may more closely mimic the “after failure” values. Not surprisingly, altitude,

cross track and angle deviations, and vertical velocity also increases “during failure”; however, these values also increase after the adaptive controller fully adapts. Given more time, these deviations and vertical velocity would settle closer to “before failure” levels.

Moreover, subjects may change their behavior while an adaptive controller is adapting with additional stick inputs. These additional stick inputs result in larger flight track errors that do eventually approach pre-failure levels given enough time. Hence, the adaptive controller should adapt as fast as possible in order to minimize flight track errors, workload, and to maintain handling qualities. This will minimize undesirable interactions between the pilot and the adaptive controller and maintain maneuvering precision.

ACKNOWLEDGMENT

The authors would like to thank John Foster and Kevin Cunningham of NASA Langley Research Center for their help in tuning the aircraft simulation model.

REFERENCES

- [1] J. T. Bosworth and P. S. Williams-Hayes, "Flight Test Results from the NF-15B Intelligent Flight Control System (IFCS) Project with Adaptation to a Simulated Stabilator Failure," presented at the AIAA Infotech@Aerospace 2007 Conference and Exhibit, Rohnert Park, CA, 2007.
- [2] A. B. Page, E. D. Meloney, and J. F. Monaco, "Flight Testing of a Retrofit Reconfigurable Control Law Architecture Using an F/A-18C," presented at the AIAA Guidance, Navigation, and Control Conference and Exhibit, Keystone, CO, 2006.
- [3] M. O. Thompson and J. R. Welsh, "Flight Test Experience with Adaptive Control System," presented at the AGARD Conference on Advanced Control Systems Concepts, 1970.
- [4] C. Wang, M. Santone, and C. Cao, "Pilot-Induced Oscillation Suppression by Using L1 Adaptive Control," *Journal of Control Science and Engineering*, vol. 2012, p. 7, 2012.
- [5] J. Hodgkinson, *Aircraft Handling Qualities*: Blackwell Science, Ltd., 1999.
- [6] A. C. Trujillo and I. M. Gregory, "Piloting Changes to Changing Aircraft Dynamics: What Do Pilots Need to Know?," presented at the 30th Digital Avionics Systems Conference, Seattle, WA, 2011.
- [7] A. C. Trujillo and I. M. Gregory, "Preliminary Exploration of Adaptive State Predictor Based Human Operator Modeling," presented at the 2012 AIAA Modeling and Simulation Technologies Conference, Minneapolis, MN, 2012.
- [8] G. E. Cooper and R. P. Harper, "The Use of Pilot Rating in the Evaluation of Aircraft Handling Qualities," AGARD April 1969.
- [9] R. P. Harper and G. E. Cooper, "Handling Qualities and Pilot Evaluation (Wright Brothers Lecture in Aeronautics)," *Journal of Guidance, Control, and Dynamics*, vol. 9, pp. 515-529, 1986 1986.
- [10] A. C. Trujillo, "Paper to Electronic Questionnaires: Effects on Structured Questionnaire Forms," presented at the HCI International 2009, San Diego, CA, 2009.
- [11] S. G. Hart and L. E. Staveland, "Development of a NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research," in *Human Mental Workload*, P. S. Hancock and N. Meshkati, Eds., ed Amsterdam: Elsevier Science Publishers B. V., 1988, pp. 139-183.